ABSTRACT OF THE DISCLOSURE

Stochastic control problems of linear systems in high dimensions are solved by modeling a structured Markov Decision Process (MDP). A state space for the MDP is a polyhedron in a Euclidean space and one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state. One or more approximations are built from above and from below to a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming.